

THE CLAIMS

What is claimed is:

1. A method for increasing failure tolerance of a storage system having a plurality of arrays, each array having a plurality of storage units, the method comprising steps of:
 - selecting a recipient array;
 - selecting a donor array from the plurality of arrays when a difference between a minimum distance of the donor array and a minimum distance of the recipient array is greater or equal to 2;
 - selecting a donor storage unit in the donor array; and
 - rebuilding at least a portion of lost recipient information from the recipient array on the selected storage unit in the donor array.
2. The method according to claim 1, wherein the minimum Hamming distance of the recipient array is $d \geq 2$ before the step of selecting the donor array from the plurality of arrays.
3. The method according to claim 1, further comprising a step of indicating to the donor array that the selected storage has been donated before the step of rebuilding the lost information on the selected storage unit.
4. The method according to claim 1, wherein the storage units are hard disk drives.
5. The method according to claim 1, wherein the storage units are RAM storage devices.

6. The method according to claim 1, further including a step of selecting a recipient storage unit from the recipient array.
7. The method according to claim 6, wherein at least a portion of the lost contents of the recipient storage unit are rebuilt onto the donor storage unit.
8. The method according to claim 1, wherein the donor storage unit is further selected based on a minimal performance impact on the donor array.
9. The method according to claim 1, further comprising a step of selecting the recipient information based on an improved performance of the recipient array.
10. The method according to claim 1, wherein the arrays of the storage system include redundancy based on an erasure or error correcting code.
11. The method according to claim 10, wherein the erasure or error correcting code is a parity code.
12. The method according to claim 10, wherein the erasure or error correcting code is a Winograd code.
13. The method according to claim 10, wherein the erasure or error correcting code is a symmetric code.
14. The method according to claim 10, wherein the erasure or error correcting

code is a Reed-Solomon code.

15. The method according to claim 10, wherein the erasure or error correcting code is an EVENODD code.

16. The method according to claim 10, wherein the erasure or error correcting code is a derivative of an EVENODD code.

17. The method according to claim 10, wherein the arrays of the storage system includes redundancy based on a product of a plurality of erasure or error correcting codes.

18. The method according to claim 17, wherein at least one of the erasure or error correcting codes is a parity code.

19. The method according to claim 17, wherein at least one of the erasure or error correcting codes is a Winograd code.

20. The method according to claim 17, wherein at least one of the erasure or error correcting code is a symmetric code.

21. The method according to claim 17, wherein at least one of the erasure or error correcting code is a Reed-Solomon code.

22. The method according to claim 17, wherein at least one of the erasure or error correcting code is an EVENODD code.

23. The method according to claim 17, wherein at least one of the erasure or error correcting code is a derivative of an EVENODD code.

24. The method according to claim 1, wherein when a storage unit in the donor array fails during the step of rebuilding at least a portion of recipient information from the recipient array on the selected storage unit, the method further comprising steps of:

terminating the step of rebuilding at least a portion of recipient information from the recipient array on the selected storage unit;

selecting a second donor array from the plurality of arrays when a difference between a minimum distance of the second donor array and a minimum distance of the second recipient array is greater or equal to 2;

selecting a donor storage unit in the second donor array; and

rebuilding at least a portion of lost recipient information from the recipient array on the selected storage unit in the second donor array.

25. The method according to claim 1, wherein when a spare storage unit becomes available, the method further comprising a step of assigning the spare storage unit to a selected array.

26. A method of increasing the failure tolerance of an array of storage units that is vulnerable to selected patterns of failures, comprising steps of:

subsequent to a storage unit failure, selecting a recipient storage unit from the array of storage units;

selecting a donor storage unit from the array of storage units such that a failure tolerance of the array is increased following a rebuild operation;

rebuilding at least a portion of lost recipient information onto the donor storage unit.

27. The method according to claim 26, wherein the minimum Hamming distance of the array is $d \geq 2$ before the step of selecting the donor storage unit.

28. The method according to claim 26, wherein the minimum Hamming distance of the array is increased upon completion of the step of rebuilding.

29. The method according to claim 26, wherein the recipient storage unit is selected based on a failure pattern of the array.

30. The method according to claim 26, wherein the donor storage unit is selected based on a predetermined target pattern.

31. The method according to claim 26, wherein the array of the storage units includes redundancy based on an erasure or error correcting code.

32. The method according to claim 31, wherein the erasure or error correcting code is a parity code.

33. The method according to claim 31, wherein the erasure or error correcting code is a Winograd code.

34. The method according to claim 31, wherein the erasure or error correcting code is a symmetric code.

35. The method according to claim 31, wherein the erasure or error correcting code is a Reed-Solomon code.

36. The method according to claim 31, wherein the erasure or error correcting code is an EVENODD code.

37. The method according to claim 31, wherein the erasure or error correcting code is a derivative of an EVENODD code.

38. The method according to claim 26, wherein the array includes redundancy based on a product of a plurality of erasure or error correcting codes.

39. The method according to claim 38, wherein at least one of the erasure or error correcting codes is a parity code.

40. The method according to claim 38, wherein at least one of the erasure or error correcting codes is a Winograd code.

41. The method according to claim 38, wherein at least one of the erasure or error correcting codes is a symmetric code.

42. The method according to claim 38, wherein at least one of the erasure or error correcting codes is a Reed-Solomon code.

43. The method according to claim 38, wherein at least one of the erasure or error

correcting codes is an EVENODD code.

44. The method according to claim 38, wherein at least one of the erasure or error correcting codes is a derivative of an EVENODD code.

45. A data storage system, comprising:
a plurality of arrays, each array having a plurality of storage units; and
a system array controller coupled to each array, the system array controller detecting a failure of a storage unit in a first array of the plurality of arrays, selecting a storage unit in a second array of the plurality of arrays when a difference between a minimum distance of the second array and a minimum distance of the first array is greater or equal to 2, and rebuilding at least a portion of information from the first array onto the selected storage unit of the second array.

46. The data storage system according to claim 45, wherein when a spare unit becomes available, the spare unit is assigned to the second array.

47. The data storage system according to claim 45, wherein at least one array has a non-uniform failure probability.

48. The data storage system according to claim 45, wherein the system array controller includes a plurality of array controllers, an array controller being coupled to at least one array of the plurality of arrays, each respective array controller detecting a failure of a storage unit in each array associated with the array controller, a first array controller of the plurality of array controllers selecting a storage unit in an array associated with the first array controller when a difference between a minimum distance of the array of the selected storage

unit and a minimum distance of an array associated with a second array controller of the plurality of array controllers is greater or equal to 2, and the first and second array controllers rebuilding at least a portion of lost information from the array associated with the second array controller onto the selected storage unit in the array associated with the first array controller.

49. The data storage system according to claim 48, wherein when a spare unit becomes available, the spare unit is assigned to the array of the selected storage unit.

50. The data storage system according to claim 48, wherein at least one array has a non-uniform failure probability.

51. The data storage system according to claim 45, wherein the arrays of the data storage system include redundancy based on an erasure or error correcting code.

52. The data storage system according to claim 51, wherein the erasure or error correcting code is a parity code.

53. The data storage system according to claim 51, wherein the erasure or error correcting code is a Winograd code.

54. The data storage system according to claim 51, wherein the erasure or error correcting code is a symmetric code.

55. The data storage system according to claim 51, wherein the erasure or error correcting code is a Reed-Solomon code.

56. The data storage system according to claim 51, wherein the erasure or error correcting code is an EVENODD code.

57. The data storage according to claim 51, wherein the erasure or error correcting code is a derivative of an EVENODD code.

58. The data storage system according to claim 45, wherein the arrays of the storage system includes redundancy based on a product of a plurality of erasure or error correcting codes.

59. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a parity code.

60. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a Winograd code.

61. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a symmetric code.

62. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a Reed-Solomon code.

63. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is an EVENODD code.

64. The data storage system according to claim 58, wherein at least one of the erasure or error correcting codes is a derivative of an EVENODD code.

65. The data storage system according to claim 45, wherein each storage unit is one of a Hard Disk Drive, a volatile Random Access Memory device, a non-volatile Random Access Memory device, an optical storage device, and a tape storage device.

66. A data storage system, comprising:
an array having a plurality of storage units, the array being vulnerable to selected patterns of failures; and
an array controller coupled to the array, the array controller detecting a failure of a first storage unit in the array, selecting a second storage unit in the array such that a failure tolerance of the array is increased following a rebuild operation, and rebuilding at least a portion of information from the first storage unit onto the second storage unit.

67. The data storage system according to claim 66, wherein the minimum Hamming distance of the array is $d \geq 2$ before the array controller selects the second storage unit.

68. The data storage system according to claim 66, wherein the minimum Hamming distance of the array is increased upon completion of rebuilding at least a portion of information from the first storage unit onto the second storage unit.

69. The data storage system according to claim 66, wherein the second storage unit is selected based on a failure pattern of the array.

70. The data storage system according to claim 66, wherein the second storage unit is selected based on a predetermined target pattern.

71. The data storage system according to claim 66, wherein the array includes redundancy based on an erasure or error correcting code.

72. The data storage system according to claim 71, wherein the erasure or error correcting code is a parity code.

73. The data storage system according to claim 71, wherein the erasure or error correcting code is a Winograd code.

74. The data storage system according to claim 71, wherein the erasure or error correcting code is a symmetric code.

75. The data storage system according to claim 71, wherein the erasure or error correcting code is a Reed-Solomon code.

76. The data storage system according to claim 71, wherein the erasure or error correcting code is an EVENODD code.

77. The data storage system according to claim 71, wherein the erasure or error correcting code is a derivative of an EVENODD code.

78. The data storage system according to claim 66, wherein the array has a non-uniform failure probability.

79. The data storage system according to claim 66, wherein each storage unit is one of a Hard Disk Drive, a volatile Random Access Memory device, a non-volatile Random Access Memory device, an optical storage device, and a tape storage device.
80. The data storage system according to claim 66, wherein the array includes redundancy based on a product of a plurality of erasure or error correcting codes.
81. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is a parity code.
82. The data storage system according to claim 80, wherein the erasure or error correcting code is a Winograd code.
83. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is a parity code.
84. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is a symmetric code.
85. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is a Reed-Solomon code.
86. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is an EVENODD code.

87. The data storage system according to claim 80, wherein at least one of the erasure or error correcting codes is a derivative of an EVENODD code.